

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
24 February 2005 (24.02.2005)

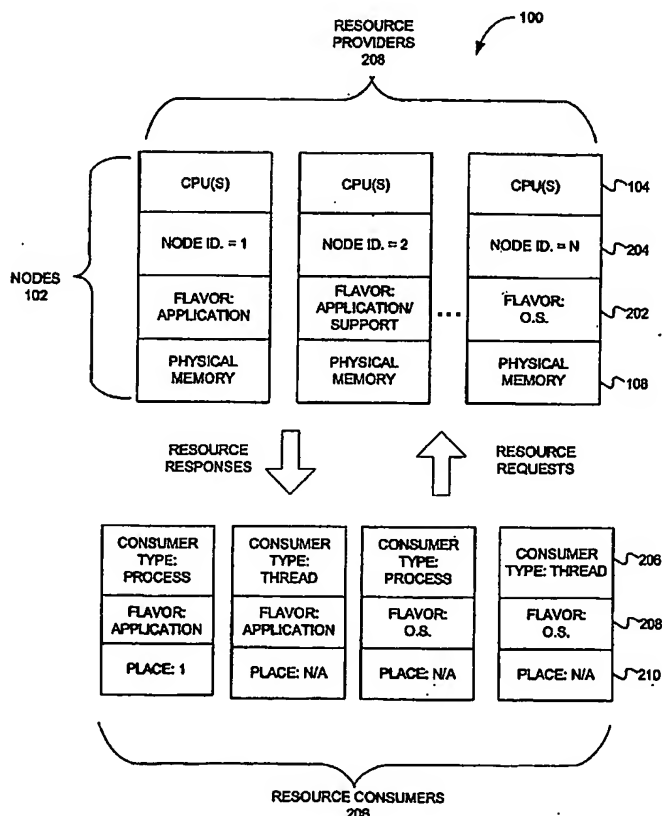
PCT

(10) International Publication Number
WO 2005/017763 A2

- (51) International Patent Classification⁷: **G06F 15/16** (74) Agents: STEFFEY, Charles, E. et al.; P.O. Box 2938, Minneapolis, Minnesota 55402 (US).
- (21) International Application Number: PCT/US2004/026675 (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (22) International Filing Date: 18 August 2004 (18.08.2004)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data: 10/643,740 18 August 2003 (18.08.2003) US
- (71) Applicant (for all designated States except US): CRAY INC. [US/US]; 411 First Avenue South, Suite 600, Seattle, Washington 98104-2860 (US).
- (72) Inventor; and
- (75) Inventor/Applicant (for US only): GIPP, Stephan Kurt [DE/US]; 10267 102nd Court West, Inver Grove Heights, Minnesota 55077 (US).
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI,

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR ALLOCATING SYSTEM RESOURCES



(57) Abstract: A system and method for allocating system resources is described herein. In one embodiment, the method comprises creating, in a computer system, a resource consumer and assigning the resource consumer one of a set of flavors. The method further includes determining whether the resource consumer is limited to receiving resources from a certain one of a set of resource providers, wherein each of the set of resource providers has one of the set of flavors. The method further includes indicating that the resource consumer is limited to receiving resources from the certain one of the set of resource providers, if the resource consumer is limited to receiving resources from the certain one of the set of resource providers. The method further includes allocating a resource to the resource consumer from one of the set of resource providers whose flavor matches the flavor assigned to the resource consumer.



SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Published:

- *without international search report and to be republished upon receipt of that report*

SYSTEM AND METHOD FOR ALLOCATING SYSTEM RESOURCES

5

LIMITED COPYRIGHT WAIVER

A portion of the disclosure of this patent document contains material to which the claim of copyright protection is made. The copyright owner has no objection to the facsimile reproduction by any person of the patent document or the patent disclosure, as it appears in the U.S. Patent and Trademark Office file or records, but reserves all other rights whatsoever.

FIELD

This invention relates generally to multiprocessor computer systems and more specifically allocating resources in multiprocessor computer systems.

15

BACKGROUND

As computer performance demands increase, multiprocessor computer systems are becoming more popular. Some popular multiprocessor architectures include nodes, which include multiple central processing units (CPUs) and multiple physical memory units. For software executing on such a multiprocessor system, the multiple physical memory units appear as one large physical memory address space. When memory is distributed over a number of nodes, system performance can vary because software processes can be assigned memory addresses on different nodes. Typically, CPUs can access local memory locations much faster than remote memory locations. Therefore, to increase system performance, it is often preferable to allocate local memory to software running on a local CPU. However, allocating remote memory may be preferable when there is not sufficient local memory available for executing high-priority software processes.

30

Some prior art multiprocessor systems provide highly flexible memory allocation schemes that allow certain physical memory segments to be allocated to specific nodes (or processors), while allowing other physical memory segments to be allocated to any node (or processor). These highly flexible memory allocation schemes often do not abstract memory allocation details from software designers, making software design relatively complicated. Moreover,

35

these highly flexible memory systems often require software to be redesigned when nodes are added or removed from the multiprocessor system.

SUMMARY

A system and method for allocating system resources is described herein.

- 5 In one embodiment, the method comprises creating, in a computer system, a resource consumer and assigning the resource consumer one of a set of flavors. The method further includes determining whether the resource consumer is limited to receiving resources from a certain one of a set of resource providers, wherein each of the set of resource providers has one of the set of flavors. The
- 10 method further includes marking a field to indicate that the resource consumer is limited to receiving resources from the certain one of the set of resource providers, if the resource consumer is limited to receiving resources from the certain one of the set of resource providers. The method further includes allocating a resource to the resource consumer from one of the set of resource
- 15 providers whose flavor matches the flavor assigned to the resource consumer.

- In one embodiment, the apparatus includes a first set of one or more nodes, wherein each of the set of nodes includes, a second set of one or more central processing units (CPUs). The apparatus further includes a physical memory communicatively coupled to each CPU of the second set, wherein the
- 20 physical memory includes a first flavor of the node, wherein the physical memory includes an operating system, and wherein the operating system is to allocate CPUs of the second set and the physical memory to resource consumers that have a second flavor that matches the first flavor.

BRIEF DESCRIPTION OF THE FIGURES

- 25 The present invention is illustrated by way of example and not limitation in the Figures of the accompanying drawings:

Figure 1 is a block diagram illustrating an exemplary computer system used in conjunction with embodiments of the invention;

- Figure 2** is a data flow diagram illustrating a flow of resource requests and responses, according to exemplary embodiments of the invention;
- 30

Figure 3 is a state machine illustrating dynamic flavor changes, according to exemplary embodiments of the invention;

Figure 4 is a flow diagram illustrating operations performed by an administrator for assigning node flavors, according to exemplary embodiments of the invention;

Figure 5 is a flow diagram illustrating operations were allocating system
5 resources, according to embodiments of the invention;

Figure 6 is a flow diagram illustrating operations for allocating memory to resource consumers based on a flavoring system, according to exemplary embodiments of the invention;

Figure 7 is a flow diagram illustrating operations for allocating CPUs to
10 resource consumers based on a flavoring system, according to exemplary embodiments of the invention; and

Figure 8 is a flow diagram illustrating operations for requesting resources, according to exemplary embodiments of the invention.

DESCRIPTION OF THE EMBODIMENTS

15 In the following description, numerous specific details are set forth. However, it is understood that embodiments of the invention may be practiced without these specific details. In other instances, well-known circuits, structures, and techniques have not been shown in detail in order not to obscure the understanding of this description. Note that in this description, references to
20 "one embodiment," "an alternative embodiment," or the like mean that the feature being referred to is included in at least one embodiment of the present invention. Further, separate references to "one embodiment" in this description do not necessarily refer to the same embodiment; however, neither are such
25 embodiments mutually exclusive, unless so stated and except as will be readily apparent to those skilled in the art. Thus, the present invention can include any variety of combinations and/or integrations of the embodiments described herein.

Herein, block diagrams illustrate exemplary embodiments of the invention. Also herein, flow diagrams illustrate operations of the exemplary
30 embodiments of the invention. The operations of the flow diagrams will be described with reference to the exemplary embodiments shown in the block diagrams. However, it should be understood that the operations of the flow diagrams could be performed by embodiments of the invention other than those

discussed with reference to the block diagrams, and embodiments discussed with references to the block diagrams could perform operations different than those discussed with reference to the flow diagrams.

Hardware and Operating Environment

5 This section provides an overview of the exemplary hardware and the operating environment in which embodiments of the invention can be practiced.

Figure 1 is a block diagram illustrating an exemplary computer system used in conjunction with embodiments of the invention. As shown in Figure 1, the multiprocessor computer system 100 includes a number of nodes 102. As
10 shown in Figure 1, each node 102 includes one or more central processing units 104 (illustrated as CPU(s) 104), hardware and/or software resources 106, and physical memory 108. In one embodiment, the hardware and/or software resources 106 include communication channels, device drivers, secondary storage, etc. Alternative embodiments call for other various hardware and/or
15 software resources.

 As shown in Figure 1, the physical memory 108 of one or more nodes 102 includes an operating system 112 (illustrated as O.S. 112). In one embodiment, the operating system 112 is software used for performing operations according to embodiments of the invention. Software is machine-
20 readable media for performing operations according to embodiments of the invention. Machine-readable media includes any mechanism that provides (i.e., stores and/or transmits) information in a form readable by a machine (e.g., a computer). For example, a machine-readable medium includes read only memory (ROM), random access memory (RAM), magnetic disk storage media,
25 optical storage media, flash memory devices, electrical, optical, and acoustical or other form of propagated signals (e.g., carrier waves, infrared signals, digital signals, etc.).

 In the computer system 100, the physical memories of all the nodes are combined to form a single logical memory address space 110. In embodiment,
30 the physical memory 108 includes random access memory and cache memory. Because the computer system 100 has a single logical memory address space, software running on one node could be allocated memory from a different node, as described in greater detail below.

Exemplary Implementation

Figure 2 is a data flow diagram illustrating a flow of resource requests and responses, according to exemplary embodiments of the invention. As shown in Figure 2, the computer system 100 includes a set of resource consumers 206 and a set of resource providers 208. As shown in Figure 2, the resource consumers 206 transmit resource requests to, and receive resource responses from, the resource providers 208. Similarly, the resource providers 208 transmit resource responses to, and receive resource requests from, the resource consumers 206. As also shown in Figure 2, each of the nodes 102 includes CPU(s) 104 and physical memory 108, as described above. Additionally, each node 102 also includes a node identifier field 204 (illustrated as NODE ID. 204) and a flavor field 202. As shown in Figure 2, each resource consumer 206 includes a consumer type field 206, a flavor field 208, and a place field 210.

In one embodiment, there are two consumer types: 1) process and 2) thread. A process is a set of instructions that resides in an allocated portion of physical memory, while a thread is an execution path (or trace) of a set of instructions, wherein the execution occurs on a CPU. In one embodiment, processes request physical memory and threads request CPUs. Alternative embodiments of the invention call for additional various consumer types such as I/O consumers, storage consumers, communication channel consumers, etc.

In one embodiment, there are three flavors including: 1) application, 2) support, and 3) O.S. Alternative embodiments, call for a greater or lesser number of flavors. In one embodiment, application-flavored resource consumers (e.g., processes and threads) are associated with application programs such as weather forecasting programs. Support-flavored are associated with shells, scripts, etc., while resource consumers of the OS-flavor are associated with operating system programs. In one embodiment, resource providers (e.g., nodes) can be assigned multiple flavors. For example, a node can be assigned both application and support flavors. However, in one embodiment, resource consumers can only be assigned one flavor. In one embodiment, resource providers can provide resources only to resource consumers that have a matching flavor. For example, an application-flavored node can only provide physical memory and CPU(s) to application-flavored resource consumers. Moreover, an

application-flavored node cannot provide physical memory and CPUs to OS and support-flavored resource consumers.

In one embodiment, the place field 210 indicates the only node 102 from which a resource consumer can acquire resources. In one embodiment, if
5 the place field 210 is not assigned a value, the place field 210 is not used in determining whether resources can be acquired from a given node, as described in greater detail below. In one embodiment, the value stored in the place field 210 corresponds to a node identifier 204 stored in one of the nodes.

Figure 3 is a state machine illustrating dynamic flavor changes,
10 according to exemplary embodiments of the invention. In one embodiment, resource consumers 206 can dynamically change flavors based on conditions within the multiprocessor computer system 100. The state machine 300 illustrates conditions under which resource consumers dynamically change flavors.

15 As shown in Figure 3, the state machine 300 is a cyclic directed graph including three nodes (illustrated as circles), where each node is a state. The three states, which correspond with resource flavors, are O.S., SUPPORT, and APPLICATION/PLACE. As shown in Figure 3, the three states are connected by a series of the directed edges (i.e., arrows), wherein each directed edge has an
20 associated condition. For example, the conditions for the two directed edges leading into the O.S. state are "system call," while the two directed edges leading out of the O.S. state are "return from system call." Moreover, the condition associated with the directed edge leading into the APPLICATION/PLACE state is "AP Run a.out." The "system call" condition
25 can occur when a system call is executed during the execution of a support-flavored or application-flavored thread. The return from system call condition can occur when an OS-flavored thread completes execution of a system call. The AP Run a.out condition can occur when a support-flavored thread executes an AP Run a.out command.

30 The state machine 300 begins at any of the states and moves to other states when the various conditions are satisfied, illustrating how resource consumers (e.g., threads) can dynamically change flavors. For example, the state changes (indicating a flavor change) from SUPPORT to O.S. when a

system call is made. As another example, the state changes from SUPPORT to APPLICATION/PLACE when an "AP Run a.out command" is executed.

While the discussion of Figures 2-3 above described the concepts of resource allocation and flavors, Figures 4-8 describe operations for assigning
5 flavors and allocating resources. In particular, Figure 4 describes a method for assigning node flavors and Figure 5 describes general operations for allocating resources based on the rules about flavor described herein. Figures 6-7 describe operations for allocating memory and processors based on the flavor rules.

Figure 4 is a flow diagram illustrating operations, performed by an
10 administrator, for assigning node flavors, according to exemplary embodiments of the invention. The flow diagram 400 will be described with reference to the exemplary computer system of Figure 1. The flow diagram 400 commences at block 402. At block 402, the first node is found. For example, using operating system software 112, an administrator goes to the first node. The process
15 continues at block 404.

At block 404, a flavor is assigned to the current node. For example, an administrator assigns a flavor to the current node. As a more specific example, an administrator uses the operating system software 112 to make a support-flavored, OS-flavored, and/or application-flavored node. As noted above, in one
20 embodiment, nodes can be assigned multiple flavors. The process continues at block 406.

As shown in block 406, it is determined whether the current node is the last node. For example, the administrator determines whether the current node is the last node to which a flavor is to be assigned. If the current node is the last
25 node, the process ends. Otherwise, the process continues at block 408.

At block 408, the next node is found. For example, using the operating system software 112, the administrator goes to the next node. From block 408, the process continues at block 404. Although the flow diagram 400 describes operations performed by an administrator, alternative embodiments call for an
30 automated process performed by the operating system software 112. For example, in one embodiment, the operating system software 112 assigns each node a default flavor at boot-up.

Figure 5 is a flow diagram illustrating operations for allocating system resources, according to embodiments of the invention. The flow diagram 500

will be described with reference to the exemplary embodiments of Figures 1 and 2. The flow diagram 500 commences at block 502, where a process and/or thread is created. For example, the operating system software 112 creates a process and/or thread. The flow continues at block 504.

5 At block 504, the process and/or thread is assigned a flavor. For example, the operating system software 112 assigns the process and/or thread a flavor. As a more specific example, the operating system software 112 marks the process' flavor field to indicate that it is a support-flavored, OS-flavored, or application-flavored process. The flow continues at block 506.

10 As shown in block 506, the place field of a process and/or thread is marked if necessary. For example, the operating system software 112 marks the place field of a process and/or thread if necessary. In one embodiment, the operating system software 112 marks the place field in processes that must achieve high performance. For example, a high-speed process stored in the
15 physical memory of node 1 should only receive memory allocations from memory local to node 1 because remote memory (i.e., memory of other nodes) would cause slow execution speeds. The flow continues at block 508.

 At block 508, an attempt is made to allocate resources to the process and/or thread. For example, the operating system software 112 attempts to
20 allocate resources to the process and/or thread. The operations of block 508 are described in greater detail below, with reference to Figures 6 - 7. From block 508, the flow ends.

 While the discussion of Figure 5 above describes general operations for creating resource consumers (e.g., threads and processes) and allocating
25 resources to those resource consumers, Figure 6-7 describe specific operations for allocating memory and CPUs to resource consumers, based on the flavor system described above.

 Figure 6 is a flow diagram illustrating operations for allocating memory to resource consumers based on a flavoring system, according to exemplary
30 embodiments of the invention. Flow diagram 600 will be described with reference to the exemplary embodiments of Figures 1-2. The flow diagram 600 commences at block 602.

At block 602, a memory allocation request is received from a process. For example, the operating system software 112 receives a memory allocation request from a process. The flow continues at block 604.

As shown in block 604, a designated node is found. For example, the
5 operating system software 112 goes to a predetermined node to begin searching for available memory. In one embodiment, the predetermined node is referred to as the affinity node. In one embodiment, the operating system software 112 begins every search at the same node, while other embodiments begin searching at different nodes based on different conditions. For example, the operating
10 system software 112 may begin searching different nodes based on the memory requestor's flavor. The flow continues at block 606.

As shown in block 606, it is determined whether the process' place field indicates that the process must receive resources from a particular node. For example, the operating system software 112 inspects the place field of the
15 process to determine whether the process must receive resources from a particular node. If the place field indicates that the process must receive resources (e.g., memory allocations) from a particular node, the flow continues at block 608. Otherwise, the flow continues at block 610.

At block 608, it is determined whether the node's flavor and ID fields
20 match the process' flavor and place fields. For example, the operating system software 112 compares the flavor and place fields of the process to the flavor and ID fields of the node. If the fields do not match, the flow ends. If the fields match, the flow continues at block 612.

At block 612, it is determined whether the node has any available
25 memory. For example, the operating system software 112 determines whether the node has any available physical memory 108. If the node has memory available, the flow continues at block 614. Otherwise, the flow continues at block 616.

At block 614, memory is allocated to the process. For example, the
30 operating system software 112 allocates a block of physical memory 108 to the process. From block 614, the flow ends.

At block 610, it is determined whether the node's flavor matches the process' flavor. For example, the operating system software 112 compares the

node's flavor field with the process's flavor field. If the fields match, the flow continues at block 612. Otherwise the flow continues at block 616.

At block 616, the next node is found. For example, the operating system software 112 goes to another node to find the needed memory. In one
5 embodiment, the operating system software 112 searches through the nodes sequentially (i.e., starting with node zero and preceding through the last node). In an alternative embodiment, the operating system software 112 finds the next node according to a search algorithm. For example, the operating system
software 112 can use the nearest node algorithm. Alternatively, the operating
10 system software 112 can perform a heuristic search for the node most likely to have the needed resources. From block 616, the process continues at block 606.

Figure 7 is a flow diagram illustrating operations for allocating CPUs to resource consumers based on a flavoring system, according to exemplary
embodiments of the invention. Flow diagram 700 will be described with
15 reference to the exemplary embodiments of Figures 1-2. The flow diagram 700 commences at block 702.

At block 702, a processor allocation request is received from a thread. For example, the operating system software 112 receives a processor allocation request from a thread. The flow continues at block 704.

20 As shown in block 704, a designated node is found. For example, the operating system software 112 goes to a predetermined node to begin searching for available CPUs. In one embodiment, the predetermined node is referred to as the affinity node. In one embodiment, the operating system software 112 begins every search at the same node, while other embodiments begin searching at
25 different nodes based on different conditions. For example, the operating system software 112 may begin searching different nodes based on the CPU requestor's flavor. The flow continues at block 706.

As shown in block 706, it is determined whether the thread's place field indicates that the thread must receive resources from a particular node. For
30 example, the operating system software 112 inspects the place field of the thread to determine whether the thread must receive resources from a particular node. If the place field indicates that the thread must receive resources (e.g., CPUs) from a particular node, the flow continues at block 708. Otherwise, the flow continues at block 710.

At block 708, it is determined whether the node's flavor and ID fields match the thread's flavor and place fields. For example, the operating system software 112 compares the flavor and place fields of the thread to the flavor and ID fields of the node. If the fields do not match, the flow ends. If the fields
5 match, the flow continues at block 712.

At block 712, it is determined whether the node has any available CPUs. For example, the operating system software 112 determines whether the node has any available CPU(s) 104. If the node has CPUs available, the flow continues at block 714. Otherwise, the flow continues at block 716.

10 At block 714, one or more CPUs is allocated to the thread. For example, the operating system software 112 allocates one or more CPU(s) 104 to the thread. From block 714, the flow ends.

At block 710, it is determined whether the node's flavor matches the thread's flavor. For example, the operating system software 112 compares the
15 node's flavor field with the thread's flavor field. If the fields match, the flow continues at block 712. Otherwise the flow continues at block 716.

At block 716, the next node is found. For example, the operating system software 112 goes to another node to find the needed CPU. From block 716, the process continues at block 706.

20 The discussion of Figures 6-7 above describes operations for allocating resources to processes and threads. In particular, Figures 6-7 describe operations for determining where to needed resources can be found and which resources can be used. In contrast, Figure 8 describes operations that resource consumers perform when requesting resources.

25 Figure 8 is a flow diagram illustrating operations for requesting resources, according to exemplary embodiments of the invention. The flow diagram 800 will be described with reference to the exemplary embodiments of Figures 1-2. The flow diagram 800 commences at block 802.

At block 802, a resource request is transmitted. For example, a resource
30 consumer 206 transmits a resource request to a resource provider 208. The process continues at block 804. At block 804, a given period of time elapses while waiting for a resource response. For example, the resource consumer 206 waits a given time period to receive a resource response. The process continues at block 806.

At block 806, it is determined whether a resource response has been received. For example, the resource consumer 206 determines whether it has received a resource response from a resource provider 208. If a response has been received, the process continues at block 810. Otherwise, the process
5 continues at block 808. As shown in block 808, the resource request is retransmitted. For example, the resource consumer 206 retransmits its resource request. The process continues at block 812.

At block 812, it is determined whether a maximum number of resource requests have been retransmitted. For example, the resource consumer 206
10 determines whether it has retransmitted the resource request a maximum number of times. If it has not retransmitted the resource request a maximum number of times, the process continues at block 804. Otherwise the process ends.

As shown in block 810, resources are accepted. For example, the resource consumer 206 accepts resources (e.g., CPUs, memory, etc.). From
15 block 810, the process ends.

Thus, a system and method for performing TLB operations have been described. Although the present invention has been described with reference to specific exemplary embodiments, it will be evident that various modifications and changes may be made to these embodiments without departing from the
20 broader spirit and scope of the invention. Accordingly, the specification and drawings are to be regarded in an illustrative rather than a restrictive sense.

CLAIMS

What is claimed is:

1. A system comprising:
a plurality of nodes, wherein each of the nodes includes,
5 a central processing units (CPU); and
a physical memory communicatively coupled to the CPU,
wherein the physical memory includes a first flavor
indicator, wherein the physical memory includes an
operating system process to allocate the CPU and the
10 physical memory to resource consumers that have a
second flavor indicator that matches the first flavor
indicator.
2. The apparatus of claim 1, wherein the resource consumers are processes
15 and threads.
3. The apparatus of claim 1, wherein the first flavor is an operating system
flavor, a support flavor, or an application flavor.
4. A method comprising:
creating, in a computer system consumer, a resource consumer;
20 assigning the resource consumer one of a set of flavors;
determining whether the resource consumer is limited to receiving
resources from a certain one of a set of resource providers,
wherein each of the set of resource providers has one of the set of
flavors;
25 if the resource consumer is limited to receiving resources from the certain
one of the set of resource providers, indicating that the resource
consumer is limited to receiving resources from the certain one of
the set of resource providers; and
allocating a resource to the resource consumer from one of the set of
30 resource providers whose flavor matches the flavor assigned to
the resource consumer.

5. The method of claim 4, wherein the field is stored in the resource consumer.
6. The method of claim 4, wherein in the resource is physical memory.
7. The method of claim 4, wherein the resource is one or more central
5 processing units.
8. The method of claim 4, wherein the set of flavors includes application, support, and operating system.
9. A machine-readable medium that provides instructions, which when executed by a machine, cause the machine to perform operations comprising:
10 creating, in a computer system consumer, a resource consumer;
assigning the resource consumer one of a set of flavors;
determining whether the resource consumer is limited to receiving
resources from a certain one of a set of resource providers,
wherein each of the set of resource providers has one of the set of
15 flavors;
if the resource consumer is limited to receiving resources from the certain
one of the set of resource providers, marking a field to indicate
that the resource consumer is limited to receiving resources from
the certain one of the set of resource providers; and
20 allocating a resource to the resource consumer from one of the set of
resource providers whose flavor matches the flavor assigned to
the resource consumer.
10. The machine-readable medium of claim 9, wherein the field is stored in the resource consumer.
- 25 11. The machine-readable medium of claim 9, wherein in the resource is physical memory.
12. The machine-readable medium of claim 9, wherein the resource is one or more central processing units.

13. The machine-readable medium of claim 9, wherein the set of flavors includes application, support, and operating system.

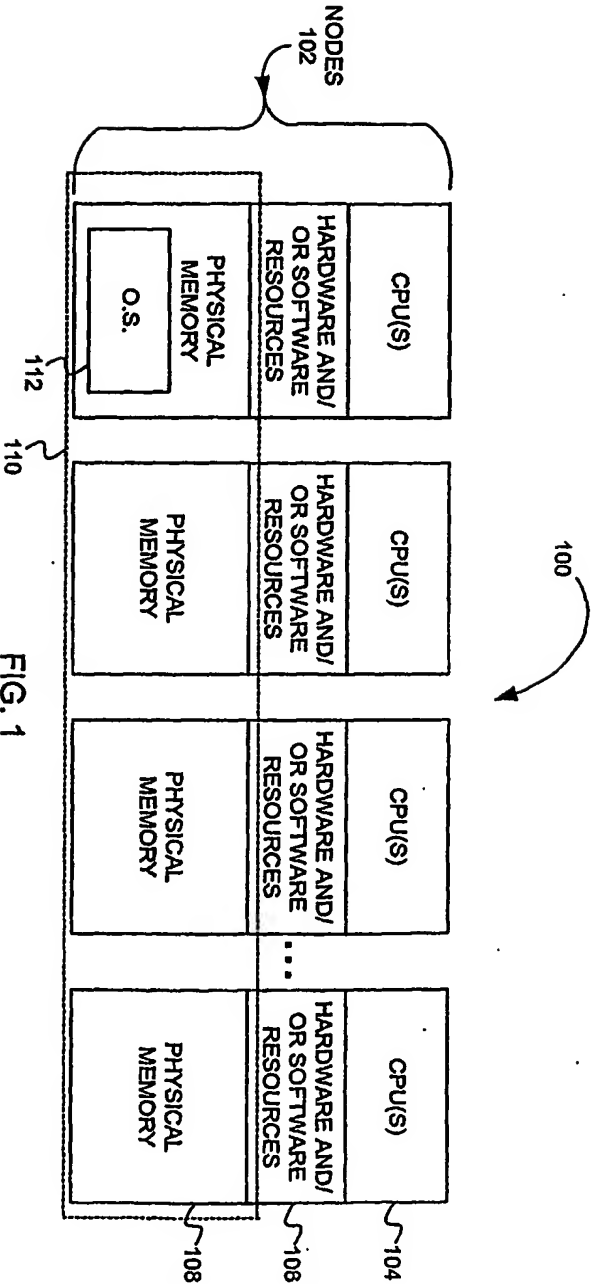


FIG. 1

2/7

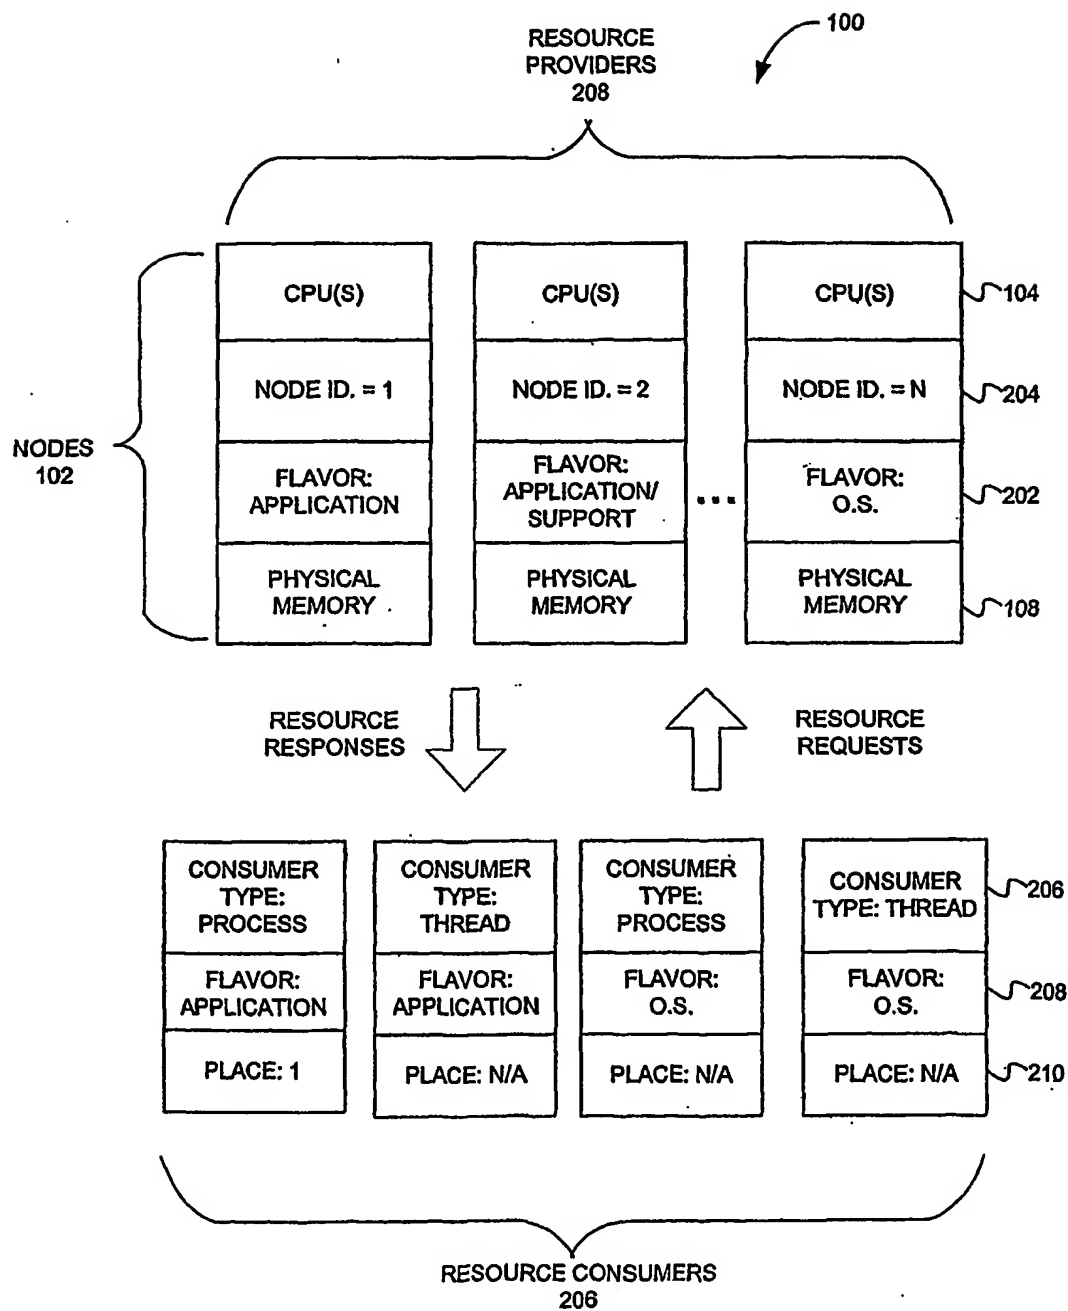


FIG. 2

3/7

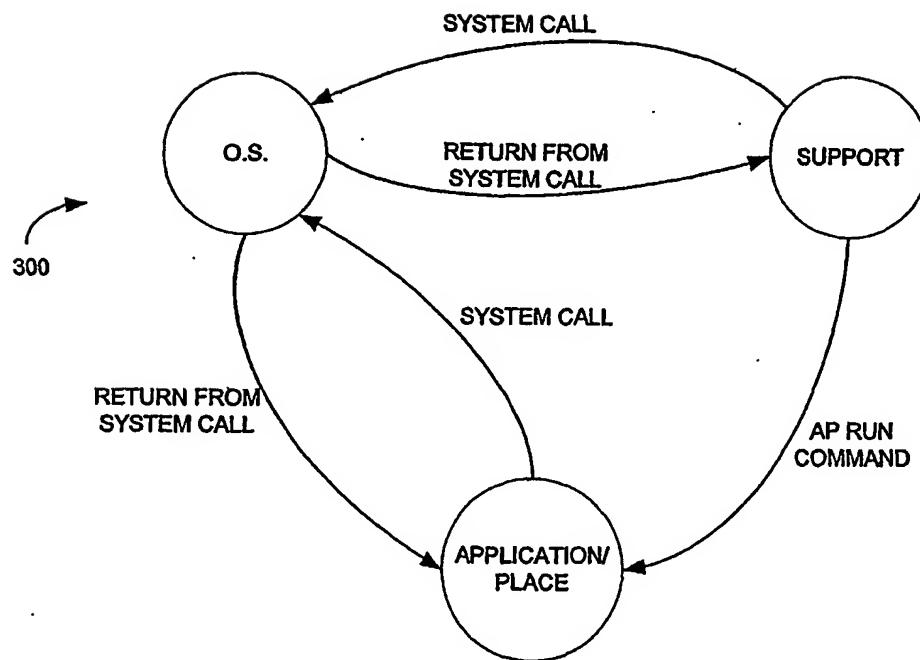


FIG. 3

4/7

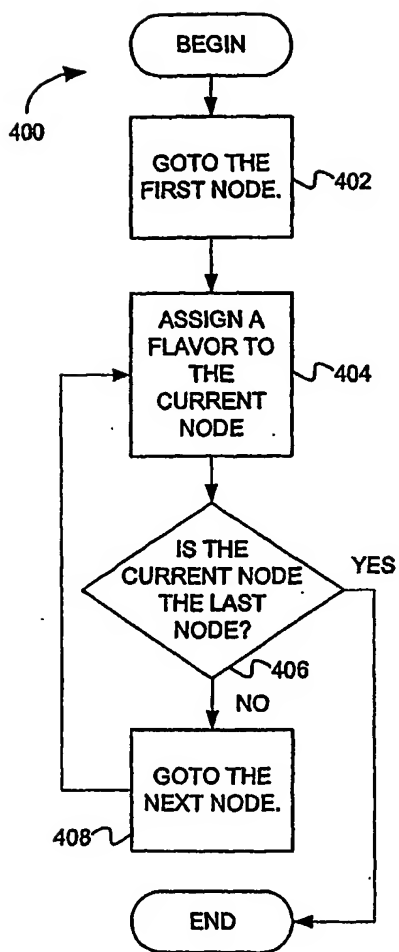


FIG. 4

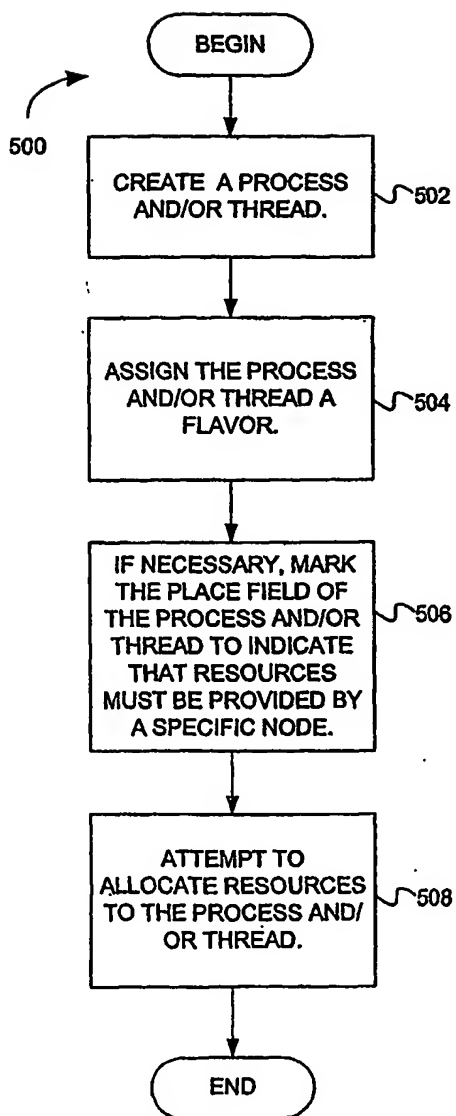


FIG. 5

5/7

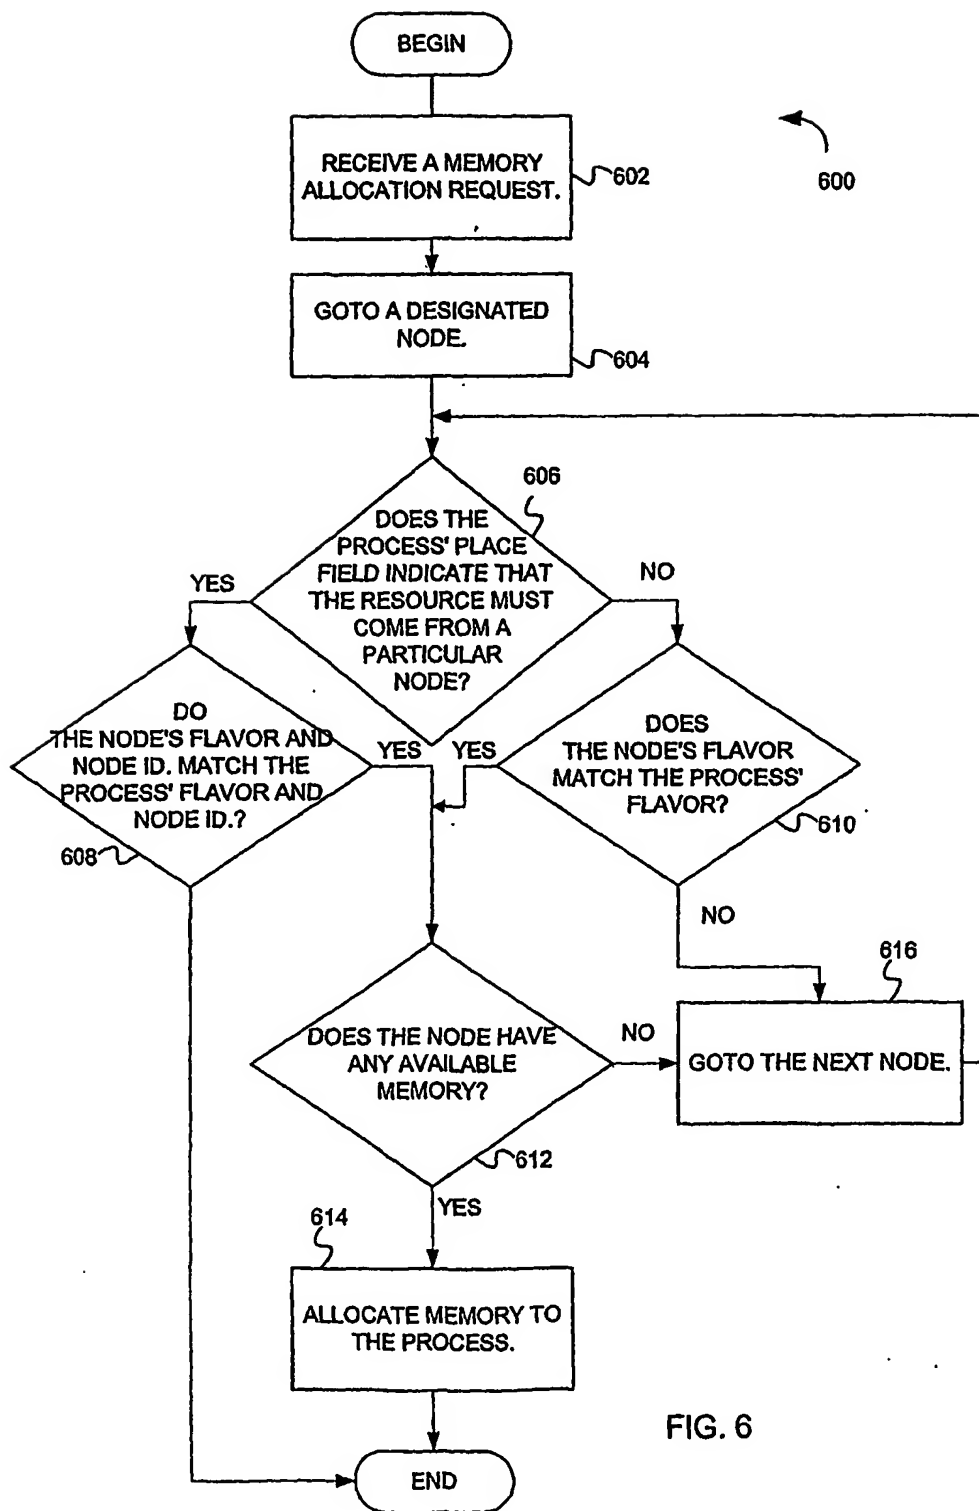


FIG. 6

6/7

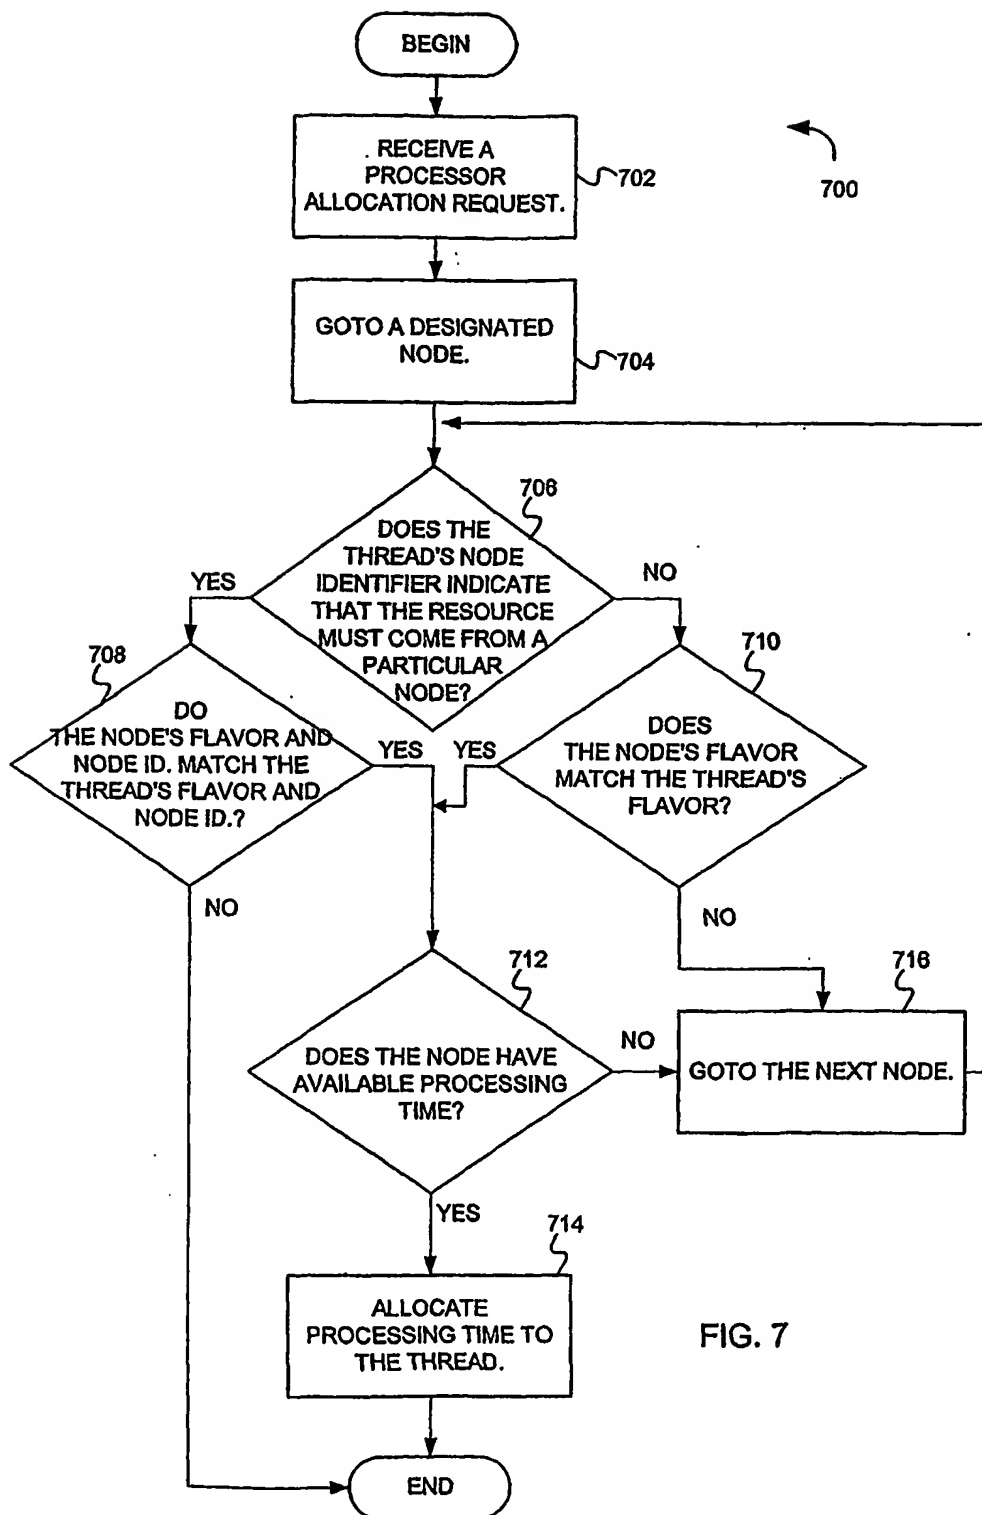


FIG. 7

7/7

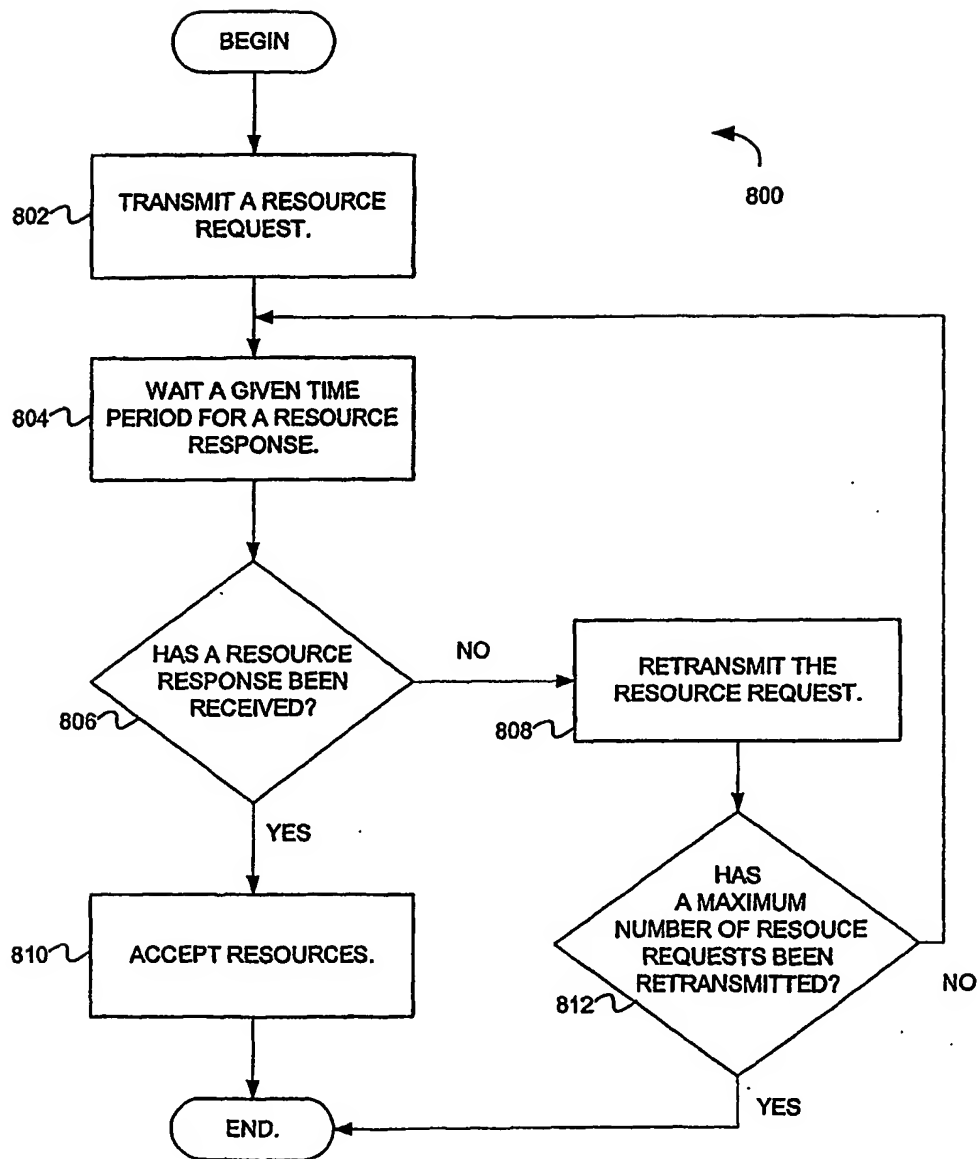


FIG. 8

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
24 February 2005 (24.02.2005)

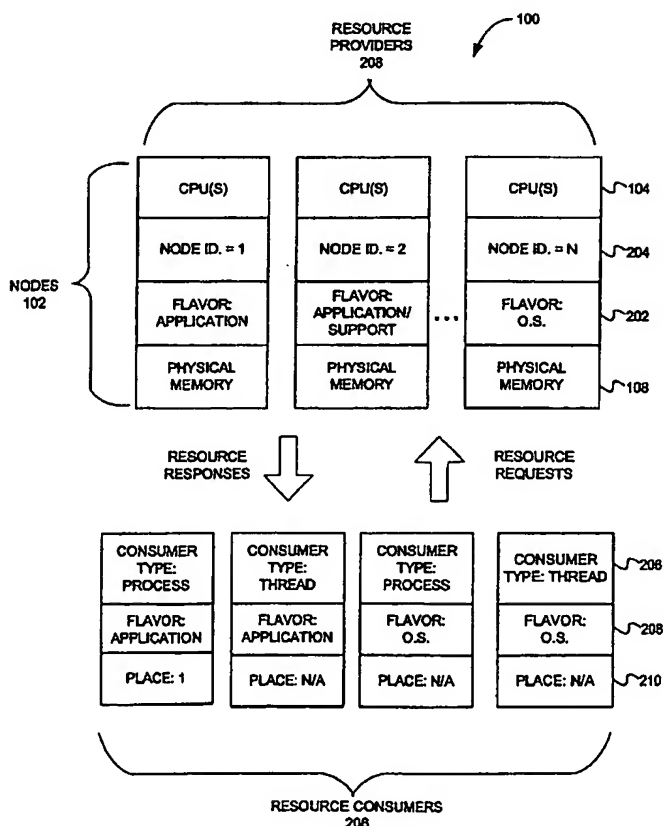
PCT

(10) International Publication Number
WO 2005/017763 A3

- (51) International Patent Classification⁷: **G06F 9/46** (74) Agents: **STEFFEY, Charles, E. et al.**; P.O. Box 2938, Minneapolis, Minnesota 55402 (US).
- (21) International Application Number: **PCT/US2004/026675** (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (22) International Filing Date: 18 August 2004 (18.08.2004)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data: 10/643,740 18 August 2003 (18.08.2003) US
- (71) Applicant (for all designated States except US): **CRAY INC.** [US/US]; 411 First Avenue South, Suite 600, Seattle, Washington 98104-2860 (US).
- (72) Inventor; and
- (75) Inventor/Applicant (for US only): **GLPP, Stephan Kurt** [DE/US]; 10267 102nd Court West, Inver Grove Heights, Minnesota 55077 (US).
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PL, PT, RO, SE, SI,

[Continued on next page]

(54) Title: SYSTEM AND METHOD FOR ALLOCATING SYSTEM RESOURCES



(57) Abstract: A system and method for allocating system resources is described herein. In one embodiment, the method comprises creating, in a computer system, a resource consumer and assigning the resource consumer one of a set of flavors. The method further includes determining whether the resource consumer is limited to receiving resources from a certain one of a set of resource providers, wherein each of the set of resource providers has one of the set of flavors. The method further includes indicating that the resource consumer is limited to receiving resources from the certain one of the set of resource providers, if the resource consumer is limited to receiving resources from the certain one of the set of resource providers. The method further includes allocating a resource to the resource consumer from one of the set of resource providers whose flavor matches the flavor assigned to the resource consumer.

WO 2005/017763 A3



SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

Published:

- *with international search report*
- *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments*

(88) Date of publication of the international search report:

20 October 2005

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US2004/026675

A. CLASSIFICATION OF SUBJECT MATTER
IPC 7 G06F9/46

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)
IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	"NAS Origin2000 Configuration For Batch Job Execution" 'Online! 20 March 2000 (2000-03-20), XP002340783 Retrieved from the Internet: URL: http://anusf.anu.edu.au/{dbs900/PBS/02K/02K-config.html} 'retrieved on 2005-10-15!	4,6-9, 11-13
A	the whole document	1-3,5,10
A	US 6 247 041 B1 (KRUEGER PHILLIP ERIC ET AL) 12 June 2001 (2001-06-12) the whole document	1-13
	----- -/-	

☒ Further documents are listed in the continuation of box C.

☒ Patent family members are listed in annex.

* Special categories of cited documents:

- *A* document defining the general state of the art which is not considered to be of particular relevance
- *E* earlier document but published on or after the international filing date
- *L* document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)
- *O* document referring to an oral disclosure, use, exhibition or other means
- *P* document published prior to the international filing date but later than the priority date claimed

- *T* later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
- *X* document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
- *Y* document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.
- *G* document member of the same patent family

Date of the actual completion of the international search

18 August 2005

Date of mailing of the international search report

30/08/2005

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Bosch Vivancos, P

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US2004/026675

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>"hptc/ClusterPack V2.0 1.8 Workload Management Tasks" 'Online! 1 July 2003 (2003-07-01), XP002340784 Retrieved from the Internet: URL: http://web.archive.org/web/20030727110012/http://www.hp.com/techservers/clusterpack_v2/1.8.html 'retrieved on 2005-10-15! the whole document</p>	1-13
A	<p>PAUL JACKSON: "Process Scheduling and Memory Placement - CpuMemSet Design Notes" 'Online! 2 July 2001 (2001-07-02), XP002340785 Retrieved from the Internet: URL: http://sourceforge.net/docman/display_doc.php?docid=8463&group_id=8875 'retrieved on 2005-10-15! the whole document</p>	1-13
A	<p>"Application Software Considerations for NUMA-Based Systems" 'Online! 5 March 2003 (2003-03-05), XP002340786 Retrieved from the Internet: URL: http://www.microsoft.com/whdc/system/platform/server/datacenter/numa_1sv.mspx 'retrieved on 2005-10-15! the whole document</p>	1-13

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US2004/026675

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 6247041	B1	12-06-2001	NONE